

Theia

Theia is located at NESCC in Fairmont, West Virginia. Theia is a Cray CS400 cluster with 1012 compute nodes with Intel Haswell processors. It has QDR InfiniBand interconnect and Panasas filesystem.

System Configuration

	Theia	Selene
CPU Type	Intel Haswell	Intel Haswell
CPU Speed	2.60 GHz	2.60 GHz
Cores/node	24	24
Total Cores	24288	528
Memory/Core	2.67 GB	2.67 GB
Memory/Node	64 GB	64 GB
Peak Flops/node	998 GFlops	998 GFlops
Total Service Nodes	10	1
Service Node Memory	256 GB	256 GB
Total BigMem Nodes	14	1
BigMem Node Memory	256 GB	256 GB
Total Flops	1010 TF	22 TF

Notes:

- Total flops is a measure of peak, and doesn't necessarily represent actual performance.
- Selene is the Test and Development System. Users must be granted specific access to the system for use.

File Systems

name	type	size	Bandwidth
scratch3	Panasas	5.5 PB	> 50 GB/s
scratch4	Panasas	5.5 PB	> 50 GB/s

Zeus

Zeus is a IBM machine located at NESCC in Fairmont, West Virginia.

System Configuration

	Zeus	Herc
CPU Type	Intel Westmere	Intel Westmere
CPU Speed	3.46 GHz	3.46 GHz
Cores/node	12	12
Total Cores	27648	576

	Zeus	Herc
Memory/Core	2 GB	2 GB
Memory/Node	24 GB	24 GB
Peak Flops/node	166 GFlops	166 GFlops
Total Service Nodes	8	1
Service Node Memory	48 GB	48 GB
Total BigMem Nodes	6	0
BigMem Node Memory	96 GB	96 GB
Total Flops	382.6 TF	8.0 TF

Notes:

- Total flops is a measure of peak, and doesn't necessarily represent actual performance.
- Herc is the Test and Development System. Users must be granted specific access to the system for use.

File Systems

name	type	size	Max Bandwidth
scratch1	Lustre	2.5 PB	22 GB/s
scratch2	Lustre	3.1 PB	28 GB/s

Gaea

Gaea is a Cray XE6 machine located in Oak Ridge, Tennessee at Oak Ridge National Laboratory (ORNL).

System Configuration

- **c1** - 368 teraflops
 - 41216 Cores
 - 32 cores/node
 - 1288 nodes
 - 157 TB of memory
 - AMD Interlagos Processors
 - 4 more login nodes that can submit to both C1 and C2
 - Initial testing begins July 31
 - Early system and user acceptance begins August 14
- **c2** - 720 teraflop Cray XE6
 - 78336 Cores
 - 32 cores/node
 - 2448 nodes
 - 157 TB of memory
 - AMD Interlagos Processors
 - 4 Login nodes (Gaea5-8)

- **t3**
 - 2994 cores
 - 184 nodes
 - AMD Interlagos 16 core
 - Nvidia K20X “Kepler” GPU
 - 2 Login nodes (Gaea9-10)
- RDTNS (remote data transfer nodes)
- LDTNS (local data transfer nodes)
- Connectivity @ 2x10GB/s to NOAA N-Wave research network

Filesystems

- 3 Files Systems
 - 1PB Fast Scratch Lustre file System (fs)
 - 4PB Staging Lustre file system (ltfs)

Jet

Jet is located at ESRL in Boulder, Colorado.

System Configuration

	vjet	njet	tjet	ujet	sjet	tgpu
CPU Type	Intel IvyBridge	Intel Nehalem	Intel Westmere	Intel Westmere	Intel SandyBridge	Intel Westmere/NVIDIA Fermi
CPU Speed	2.6 GHz	2.8 GHz	2.66 GHz	2.66 GHz	2.6 GHz	2.8 GHz
Cores/node	16	8	12	12	16	12
Total Cores	4608	3584	10128	6048	5440	132
Memory/Core	4GB	3 GB	2 GB	2 GB	2GB	2GB
Memory/Node	64GB	24 GB	24 GB	24 GB	32 GB	24 GB
Available Mem/Node*	60.9GB	22.1GB	22.1GB	22.1GB	29.8GB	22.1GB
Peak Flops/node	332.8 GF	89.6 GF	127.7 GF	127.7 GF	332.8 GF	***
Relative Perf/Core	1.65	1.18	1.00	1.00	1.44	1.00
Interconnect	FDR Infiniband	QDR Infiniband	QDR Infiniband	QDR Infiniband	QDR Infiniband	QDR Infiniband
Total Flops	93.6 TF	40.1 TF	107.7 TF	64.4 TF	113.2 TF	***

Notes:

- Total flops is a measure of peak, and doesn't necessarily represent actual performance.
- Relative performance is based on an average of FIM, WRF, and GFS run on 3 different node counts for each code. The relative performance is the average performance of the different models run at different node counts.
- Although tjet performance is slower per core than njet, each tjet node is approximately 20%

faster than an njet node (12 vs. 8 cores)

- Available Mem/Node is the total memory available to application. The difference between this value and the total available memory is due to OS overhead and other system buffers.

tjet and ujet are the same architecture, just purchased at different times. The schedule is transparent and jobs will run on either system

- The tgpu system is a small testbed based on tjet nodes, but each node contains two NVIDIA Fermi M2050 GPU. Each GPU has 3GB of RAM. The CUDA programming environment is available as a loadable module on the system. If you have more questions about accessing this system, please send a message to the [Help System](#).

File Systems

name	type	size
lfs2	Lustre	1100 TB
lfs3	Lustre	1.6 PB
pan2	Panasas	653 TB

From:
<https://rdhpcs-common-docs.rdhpcs.noaa.gov/wikis/rdhpcs-common-docs/> - **RDHPCS-Common-Docs**

Permanent link:
https://rdhpcs-common-docs.rdhpcs.noaa.gov/wikis/rdhpcs-common-docs/doku.php?id=program_information&rev=1435981959 

Last update: **2015/07/04 04:52**